



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/658,871	09/09/2003	Rikard M. Kjellberg	3399P110CX	5159

26529 7590 07/13/2007
BLAKELY SOKOLOFF TAYLOR & ZAFMAN/PDC
1279 OAKMEAD PARKWAY
SUNNYVALE, CA 94085-4040

EXAMINER

ASSESSOR, BRIAN J

ART UNIT	PAPER NUMBER
----------	--------------

2114

MAIL DATE	DELIVERY MODE
-----------	---------------

07/13/2007

PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Office Action Summary

Application No.

10/658,871

Applicant(s)

KJELLBERG, RIKARD M.

Examiner

Brian J. Assessor

Art Unit

2114

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 19 April 2007.
- 2a) ☒ This action is **FINAL**. 2b) ☐ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-21, 23 and 24 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-21, 23 and 24 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☒ The drawing(s) filed on 09 September 2003 is/are: a) ☒ accepted or b) ☐ objected to by the Examiner.
- Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
- Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
2. ☐ Certified copies of the priority documents have been received in Application No. _____.
3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).
- * See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- 1) ☐ Notice of References Cited (PTO-892)
- 2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
- 3) ☐ Information Disclosure Statement(s) (PTO/SB/08)
Paper No(s)/Mail Date _____
- 4) ☐ Interview Summary (PTO-413)
Paper No(s)/Mail Date. _____
- 5) ☐ Notice of Informal Patent Application
- 6) ☐ Other: _____

DETAILED ACTION

Claims 1-21, 23, and 24 are pending. Claims 1, 2, 7, 9, 11, 13, 15, 17, 19, and 24 have been amended and are addressed below.

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

Claims 1-21, 23, and 24 are rejected under 35 U.S.C. 103(a) as being unpatentable over Brown (4,710,926) in view of Li (5,473,599) in further view of Garg (20050193229).

As per claim 1, Brown teaches:

A method comprising:

maintaining a plurality of processes of a particular service type in a processing system; (Brown column 3, lines 56-59; each processor runs a process and all are the same service type.)

assigning a status to each of the processes; (Brown column 4, lines 55-65)

causing each of the processes to monitor the other processes of said service type; (Brown column 7, lines 5-11; each processor maintains a list of the states of each other processor which is monitoring.)

Brown does not explicitly disclose a method for causing said plurality of processes to interact with each other to establish a priority of status, such that each of said plurality of processes can alter the priority of another of said plurality of processes without the use of a master to enable said interaction or alteration of priority. Brown also does not explicitly disclose a method wherein said priority is based on a value of the unique process identifier assigned to each of said plurality of processes.

In column 2, lines 44-64; column 3, lines 5-14; column 9, lines 34-42 and in column 9, lines 29-31 respectively Li discloses a method in which multiple processes determine the priority status of each of the processes. It would have been obvious to a person of ordinary skill in the art at the time of invention to include the priority status method as taught by Li, in order to create a more efficient stand-by takeover system. This would have been obvious because Li clearly teaches that the above system is better suited for a more simplistic device failover within a network. (Li column 2, lines 9-12)

Brown and Li fail to explicitly disclose a method for assigning a unique process identifier to each of the plurality of processes.

On page 2, paragraphs 0020-0022, Garg clearly discloses a method wherein within a distributed fault tolerant computer system needs to use unique process identifiers. It would have been obvious to a person of ordinary skill in the art to include the unique identifier method as taught by Garg in order to create a better managed fault tolerant distributed system. This would have been obvious because Garg teaches that

Art Unit: 2114

the above method is better suited for handling faults in a distributed system. (Garg page 1, paragraphs 0005-0006)

As per claim 2, Brown teaches:

A method as recited in claim 1, further comprising:

causing each of the processes to respond autonomously to a predetermined condition by changing its own status between active and non-active. (Brown column 7, lines 19-31; when a processor fails the other processors check to see if they are the spare processor for the failed process and if it is then it takes over for the failed processor.) wherein the predetermined condition involves another process of the particular service type. (Brown column 7, lines 27-31; the processes check their SpareGroup table to determine if they should take over for a failed process.)

As per claim 3, Brown teaches:

A method as recited in claim 2, further comprising causing each of the processes independently to maintain a list of other participant processes in the processing system. (Brown column 7, lines 5-11; each processor keeps track of heartbeats received/not received.)

As per claim 4, Brown teaches:

A method as recited in claim 3, wherein the plurality of processes includes an active process and a non-active process corresponding to the active process, each

independently maintaining said list. (Brown column 7, lines 5-11; each processor keeps track of heartbeats received/not received or active/non-active.)

As per claim 5, Brown teaches:

A method as recited in claim 4, wherein the non-active process can autonomously change its status to active in response to an event affecting the active process. (Brown column 7, lines 23-31; each process has a list of which processors it takes over for, and can be in a non-active state when this occurs. Figure 5, from stand-by to active)

As per claim 6, Brown teaches:

A method as recited in claim 1, further comprising:

causing each of the processes to send heartbeat messages to each other process; (Brown column 7, lines 3-5)

causing each of the processes to listen for heartbeat messages from other processes; (Brown column 7, lines 9-11)

causing each of the processes to update its list of participant processes based on receipt of heartbeat messages from other processes; (Brown column 7, lines 9-11)

causing each of the processes to update its list of participant processes based on the lack of receipt of heartbeat messages from other processes from which heartbeat messages have previously been received. (Brown column 7, lines 19-23; the table is updated once the Pulse Count is reduced to zero.)

As per claim 7, Brown teaches:

A method as recited in claim 1, wherein each process determines its status based on its unique process identifier. (Brown column 4, lines 43-45)

As per claim 8, Brown teaches:

A method as recited in claim 7, wherein each process determines its status based on the value of its unique process identifier relative to the value of the unique identifier of each other process. (Brown column 4, lines 43-45; with respect to arrangement 100)

As per claim 9, Brown teaches:

A method comprising:

introducing a plurality of processes into a processing system, each of the processes having a service type; (Brown column 3, lines 56-59; each processor runs a process and all are the same service type.)

assigning a status to each of the processes, each said status selected from among a plurality of prioritized statuses, (Brown column 4, lines 55-56) including a primary status (Brown column 4, lines 56-57) and a standby status, (Brown column 4, lines 57-59) such that at least one of the processes is a primary process and at least one of the processes is a standby process for the primary process; (Brown column 4, lines 55-59)

Art Unit: 2114

maintaining each of the processes so that each of the processes monitors its own status and the status of each other process of the same service type (Brown column 7, lines 3-11; each processor listens for others heartbeats and/or lack of heartbeats.) and can change its status from standby to primary without the use of a master, in response to an external event relating to a process of said same service type. (Brown column 7, lines 27-31; figure 5 shows the process from stand-by to active is done automatically.)

Brown does not explicitly disclose a method for causing said plurality of processes to interact with each other to establish a priority of status, such that each of said plurality of processes can alter the priority of another of said plurality of processes without the use of a master to enable said interaction or alteration of priority. Brown also does not explicitly disclose a method wherein said priority is based on a value of the unique process identifier assigned to each of said plurality of processes.

In column 2, lines 44-64; column 3, lines 5-14; column 9, lines 34-42 and in column 9, lines 29-31 respectively Li discloses a method in which multiple processes determine the priority status of each of the processes. It would have been obvious to a person of ordinary skill in the art at the time of invention to include the priority status method as taught by Li, in order to create a more efficient stand-by takeover system. This would have been obvious because Li clearly teaches that the above system is better suited for a more simplistic device failover within a network. (Li column 2, lines 9-12)

Brown and Li fail to explicitly disclose a method for assigning a unique process identifier to each of the plurality of processes.

On page 2, paragraphs 0020-0022, Garg clearly discloses a method wherein within a distributed fault tolerant computer system needs to use unique process identifiers. It would have been obvious to a person of ordinary skill in the art to include the unique identifier method as taught by Garg in order to create a better managed fault tolerant distributed system. This would have been obvious because Garg teaches that the above method is better suited for handling faults in a distributed system. (Garg page 1, paragraphs 0005-0006)

As per claim 10, Brown teaches:

A method as recited in claim 9, further comprising causing each of the processes to maintain a list of other participant processes in the processing system. (Brown column 7, lines 5-11; each processor keeps track of heartbeats received/not received.)

As per claim 11, Brown teaches:

A method as recited in claim 9, wherein each process determines its status based on its unique process identifier. (Brown column 4, lines 43-45)

As per claim 12, Brown teaches:

A method as recited in claim 11, wherein each process determines its status based on the value of its unique process identifier relative to the value of the unique

identifier of each other process of the same service type. (Brown column 4, lines 43-45; with respect to arrangement 100)

As per claim 13, Brown teaches:

A method comprising:

introducing a plurality of processes into a processing system, each process having a service type; (Brown column 3, lines 56-59; each processor runs a process and all are the same service type.)

causing each of the processes independently to maintain a list of other participant processes in the processing system. (Brown column 7, lines 9-11)

causing each of the processes to send a heartbeat message repeatedly to each other process; (Brown column 7, lines 3-5)

causing each of the processes to listen for heartbeat messages from other processes; (Brown column 7, lines 9-11)

causing each of the processes to update its list of participant processes based on receipt of heartbeat messages from other processes; (Brown column 7, lines 9-11)

causing each of the processes to update its list of participant processes based on the lack of receipt of heartbeat messages from other processes from which heartbeat messages have previously been received; (Brown column 7, lines 19-23; once the Pulse Count is decremented to zero the table will be updated.)

enabling each of the processes to select a status for itself, from among a plurality of prioritized statuses, including a primary and a standby status, without the use of a

master, such that the plurality of processes includes a primary process and a standby process for the primary process. (Brown column 4, lines 55-65; each active processor has a processor to act as a stand-by, and the transition from stand-by to active, as can be seen in figure 5, is automatic.)

Brown does not explicitly disclose a method for causing said plurality of processes to interact with each other to establish a priority of status, such that each of said plurality of processes can alter the priority of another of said plurality of processes without the use of a master to enable said interaction or alteration of priority. Brown also does not explicitly disclose a method wherein said priority is based on a value of the unique process identifier assigned to each of said plurality of processes.

In column 2, lines 44-64; column 3, lines 5-14; column 9, lines 34-42 and in column 9, lines 29-31 respectively Li discloses a method in which multiple processes determine the priority status of each of the processes. It would have been obvious to a person of ordinary skill in the art at the time of invention to include the priority status method as taught by Li, in order to create a more efficient stand-by takeover system. This would have been obvious because Li clearly teaches that the above system is better suited for a more simplistic device failover within a network. (Li column 2, lines 9-12)

Brown and Li fail to explicitly disclose a method for assigning a unique process identifier to each of the plurality of processes.

On page 2, paragraphs 0020-0022, Garg clearly discloses a method wherein within a distributed fault tolerant computer system needs to use unique process

Art Unit: 2114

identifiers. It would have been obvious to a person of ordinary skill in the art to include the unique identifier method as taught by Garg in order to create a better managed fault tolerant distributed system. This would have been obvious because Garg teaches that the above method is better suited for handling faults in a distributed system. (Garg page 1, paragraphs 0005-0006)

As per claim 14, Brown teaches:

A method as recited in claim 13, wherein for each process, the selection of status is based on the value of the unique process identifier of said process relative to the value of the unique process identifier of other processes having the same service type as said process. (Brown column 4, lines 43-45; with respect to arrangement 100)

As per claim 15, Brown teaches:

A processing system comprising:

a plurality of processes, each process having a service type; (Brown column 3, lines 56-59; each processor runs a process and all are the same service type.)

means for assigning a status to each of the processes, each said status selected from among a plurality of prioritized statuses, (Brown column 4, lines 55-56) including an active status (Brown column 4, lines 56-57) and a standby status, Brown column 4, lines 57-59) such that at least one of the processes is a primary process and at least one of the processes is a standby process for the primary process; (Brown column 4, lines 55-59)

means for maintaining each of the processes so that each of the processes monitors its own status and the status of each other process of the same service type (Brown column 7, lines 3-11; each processor listens for others heartbeats and/or lack of heartbeats.) and can autonomously change its status from standby to primary in response to an external event. (Brown column 7, lines 27-31; figure 5 shows the process from stand-by to active is done automatically.)

Brown does not explicitly disclose a method for causing said plurality of processes to interact with each other to establish a priority of status, such that each of said plurality of processes can alter the priority of another of said plurality of processes without the use of a master to enable said interaction or alteration of priority. Brown also does not explicitly disclose a method wherein said priority is based on a value of the unique process identifier assigned to each of said plurality of processes.

In column 2, lines 44-64; column 3, lines 5-14; column 9, lines 34-42 and in column 9, lines 29-31 respectively Li discloses a method in which multiple processes determine the priority status of each of the processes. It would have been obvious to a person of ordinary skill in the art at the time of invention to include the priority status method as taught by Li, in order to create a more efficient stand-by takeover system. This would have been obvious because Li clearly teaches that the above system is better suited for a more simplistic device failover within a network. (Li column 2, lines 9-12)

Brown and Li fail to explicitly disclose a method for assigning a unique process identifier to each of the plurality of processes.

On page 2, paragraphs 0020-0022, Garg clearly discloses a method wherein within a distributed fault tolerant computer system needs to use unique process identifiers. It would have been obvious to a person of ordinary skill in the art to include the unique identifier method as taught by Garg in order to create a better managed fault tolerant distributed system. This would have been obvious because Garg teaches that the above method is better

As per claim 16, Brown teaches:

A processing system as recited in claim 15, further comprising means for causing each of the processes to maintain a list of other participant processes in the processing system. (Brown column 7, lines 5-11; each processor keeps track of heartbeats received/not received.)

As per claim 17, Brown teaches:

A processing system as recited in claim 15, wherein each process determines its status based on its unique process identifier. (Brown column 4, lines 43-45)

As per claim 18, Brown teaches:

A processing system as recited in claim 15, wherein each process determines its status based on the value of its unique process identifier relative to the value of the unique identifier of each other process of the same service type. (Brown column 4, lines 43-45; with respect to arrangement 100)

As per claim 19, Brown teaches:

A method comprising:

maintaining a plurality of processes in a processing system, each process having an ability to independently monitor a status of each other process of said plurality of processes, without the use of a master; (Brown column 7, lines 5-11)

Brown fails to explicitly disclose a method causing said plurality of processes to interact with each other to establish a priority of status, such that each of said plurality of processes can alter the priority of another of said plurality of processes without the use of a master to enable said interaction or alteration of priority and wherein said priority is based on the unique process value of an identifier assigned to each of said plurality of processes..

In column 2, lines 44-64; column 3, lines 5-14; column 9, lines 34-42; and in column 9, lines 29-31 respectively Li teaches a method for establishing a priority status of devices throughout a system. It would have been obvious to a person of ordinary skill in the art at the time of invention to include the priority status method as taught by Li, in order to create a more efficient stand-by takeover system. This would have been obvious because Li clearly teaches that the above system is better suited for a more simplistic device failover within a network. (Li column 2, lines 9-12)

Brown and Li fail to explicitly disclose a method for assigning a unique process identifier to each of the plurality of processes.

Art Unit: 2114

On page 2, paragraphs 0020-0022, Garg clearly discloses a method wherein within a distributed fault tolerant computer system needs to use unique process identifiers. It would have been obvious to a person of ordinary skill in the art to include the unique identifier method as taught by Garg in order to create a better managed fault tolerant distributed system. This would have been obvious because Garg teaches that the above method is better

As per claim 20:

A method as recited in claim 19, wherein said interaction and said alteration amongst said plurality of processes is used to enable fault tolerance for at least one of said processes in said processing system. (Li column 2, lines 15-40; shows an organized system for setting the primary device on a network, through means of priority.)

As per claim 21:

A method as recited in claim 19, wherein said status is one of: primary, to become primary, or standby. (Li column 2, lines 44-46; a higher priority router takes over the primary status of the existing lower priority router.)

As per claim 23:

A method as recited in claim 19, wherein said priority is further based on the status assigned to each of said plurality of processes. (Li column 3, lines 2-5)

As per claim 24, Brown teaches:

A method for providing fault tolerance in a processing system, the method comprising:

enabling a plurality of processes in a processing system each to broadcast a periodic heart-beat message, wherein said heart-beat message includes an identifier for each of said plurality of processes; (Brown column 7, lines 5-11)

enabling each of said plurality of processes to receive each said heart-beat message; (Brown column 7, lines 5-11)

causing each of said plurality of processes to maintain an individual record of said plurality of processes; (Brown column 7, lines 5-11)

causing each of said plurality of processes to update said individual record based on said heart-beat messages; (Brown column 7, lines 5-11)

assigning each of said processes with a status, wherein said status is one of: primary, to become primary, or standby; (Brown column 4, lines 55-59)

enabling said plurality of processes to negotiate a hierarchy of control amongst each other based on the broadcast and receipt of heart-beat messages by each of said plurality of processes, (Brown column 7, lines 3-5)

Brown does not explicitly disclose a method for causing said plurality of processes to interact with each other to establish a priority of status, such that each of said plurality of processes can alter the priority of another of said plurality of processes without the use of a master to enable said interaction or alteration of priority. Brown

also does not explicitly disclose a method wherein said priority is based on a value of the unique process identifier assigned to each of said plurality of processes.

In column 2, lines 44-64; column 3, lines 5-14; column 9, lines 34-42 and in column 9, lines 29-31 respectively Li discloses a method in which multiple processes determine the priority status of each of the processes. It would have been obvious to a person of ordinary skill in the art at the time of invention to include the priority status method as taught by Li, in order to create a more efficient stand-by takeover system. This would have been obvious because Li clearly teaches that the above system is better suited for a more simplistic device failover within a network. (Li column 2, lines 9-12)

Brown and Li fail to explicitly disclose a method for assigning a unique process identifier to each of the plurality of processes.

On page 2, paragraphs 0020-0022, Garg clearly discloses a method wherein within a distributed fault tolerant computer system needs to use unique process identifiers. It would have been obvious to a person of ordinary skill in the art to include the unique identifier method as taught by Garg in order to create a better managed fault tolerant distributed system. This would have been obvious because Garg teaches that the above method is better

Response to Arguments

Applicant's arguments filed 4/19/07 have been fully considered but they are not persuasive.

Applicant's Argument:

Applicant argues that the combination of Brown, Li, and Garg does not teach a priority based on a unique process identifier of each process, and further argues that there is no motivation to combine them.

Examiner's Response:

The purpose of using Garg is just to display that using unique identifiers is well known in the art, official notice could just as easily be used. Unique identifiers have been used in many cases from process identifiers, to serial numbers, and IP addresses used to identify systems over a network. The combination is just an example of the many already known applications of unique identifiers. The remaining limitation of establishing priorities based on the process identifiers is taught entirely by Li. Therefore, examiner respectfully maintains the rejection.

Applicant's Argument:

Applicant argues that Brown does not disclose determining a status of the processor, wherein the status is based on the value of its process identifier and the value of the other process identifiers.

Examiner's Response:

Brown teaches having has a physical ID and a logical ID as cited in the prior action. These IDs are combined in order to form a single unique ID, which the logical ID is used to determine the status and purpose of the process. It is additionally pointed out that Li also uses the process IDs in order to determine the primary and secondary process' which also teaches the use of the process identifier in order to determine a status. Therefore, the examiner respectfully maintains the rejection.

Conclusion

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Brian J. Assessor whose telephone number is (571) 272-0825. The examiner can normally be reached on M-F 9:30-6:30.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Scott Baderman can be reached on (571)272-3644. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Art Unit: 2114

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

BA



GABRIEL CHU
PRIMARY EXAMINER